Oral History of Ken Haughton

Interviewed by:
Jim Porter

Recorded: February 14, 2007
Mountain View, California

CHM Reference number: X3888.2007

Jim Porter: We're here today to have a discussion with Dr. Kenneth E. Haughton, who has a very interesting background and history of experiences with some of the more important events that occurred in the disk drive industry. Ken, I wonder if we could start by getting a little of your history as to how you got started as a person, with your education, and your early experiences.

Ken Haughton: OK. I'm not sure how far back you want to go, but I can go quite a ways. I was raised by a family in Oregon, and my father was a logger. And in fact, when I was in high school, I worked in the logging camps for some time, and in 1946 I decided that I didn't want to spend the rest of my life in the logging camps, and that I should go to college. My family couldn't afford to send me, so at that time we still had the World War II G.I Bill of Rights, so I enlisted in the Marine Corps. Went into the Marine Corps for two years, and then came -- and that brought me to California. And in the fall of 1948 I started college at the San Francisco City College -- community college in San Francisco. I went there for two years, and after two years I transferred to Berkeley at the University of California, and was there for two more years, and graduated in 1952. At that time I took a job with the General Electric Company at Hanford, Washington. That was a design job. I was designing machinery to machine graphite, because they were making a couple of new reactors at that particular point in time, and we needed the equipment to cut that graphite. Interesting little side point there is that there was a jurisdictional dispute going on amongst the Kaiser Engineers then -- is graphite more like wood or is it more like metal? Because the millwrights will do it if it's like wood, and the machinists will do it if it's like metal. So we had a very interesting jurisdictional dispute there. Turns out everyone was qualified with the union card of whichever one won, so they just had a nice face-off and eventually -- I don't even remember who won. From General Electric I went to Iowa State College, and at Iowa State I took a position there teaching engineering graphics and doing part-time graduate work.

Porter: Let me just ask the question, Ken. How did you transition from California to Iowa? What made you go there?

Haughton: Oh, well, because they were looking for a faculty member.

Porter: Ah, a job.

Haughton: A job. A job took me to Iowa. So I went back there, and for three years I taught at Iowa State. And as I finished my Masters Degree I was recruited by a professor at Cornell University to come to Cornell and work on a PhD. I agreed to do that. In the fall of -- let's see, that's 1956 -- I went to Cornell and, lo and behold, the guy had taken a job and wasn't there -- the faculty member that had recruited me. So I taught in the Mechanics Department at Cornell University for one year, and left, and joined IBM in

San Jose, and came back to California after that one year. There's a lot less snow in California than there is in New York.

Porter: So you went to work for IBM and found that doing a lot of interesting things on some kind of a storage products…?

Haughton: That's right. I came to work on July 1st 1957, which is close to -- it's the exact day they moved into the research building which became known as Building 25 on the IBM campus in San Jose. And it's very close to the day when they really started shipping the RAMAC. So they had a couple of major things happening then. At the time I arrived, one of the biggest problems with the RAMAC, believe it or not, was that the compressor that supplied the air for the hydrostatic air bearings to hold the heads off the disk was the least reliable thing in the system. They had a box about -- a cube about 3 feet on the side…

Porter: And it was very noisy.

Haughton: Very noisy, and it broke frequently. And so somebody had the wild idea of: We ought to be able to design some heads that use the ambient air instead of supplying compressed air, and I was assigned to work on that project when I arrived at the -- I came into IBM in the Research Division, and this was going on in the Product Division, and so I was loaned from the Research Division to the Product Division to do this work on the hydrodynamic bearings. There were three of us to start with that were working together to develop these bearings. Shortly one of them was moved on, and two of us worked together. It just happened by coincidence that the other person besides myself was a person that I went to college with at Berkeley, and we had worked together before, so we got along very, very well, and spent some time. There are a few interesting things that happened, a lot of interesting things for that matter, that happened on that, but I'll tell you one example. As we manufactured, and we tended to manufacture them ourselves -- we had a lapping machine, and we had a machine shop behind us if there was needed to be some machine shop work done. As we made these heads, we discovered that some of them flew very, very well. They were -- they had a stiff spring constant, if I can call it that. The deflection per pound of change was such that it took quite -- carried quite a big load. But some of them wouldn't work at all. So, we said "Well, maybe we aren't getting them flat enough." And so we worked very hard at developing a way to flatten -- make the bearing flat. And by the way, I almost threw my glasses away when I did that. We went to an optician that manufactured glasses and said "We need to make this flat head. How do you do it?" So he showed us what he did, and I was convinced that -- I couldn't understand how I could see though my glasses. They were so crude compared to what we were trying to do. Anyway, we finally managed a technique to get them flat, very flat. And, lo and behold, they all failed.

So there's an example of where failure can teach you something. And in fact, we immediately set about to figure out how to make them cylindrical so that they would go into the direction of motion presenting a curved surface to the air that was there. That was a little bit of a trick in itself, because these bearings were about three-quarters of an inch long and they were a little over a half an inch wide, and we needed a hump on the cylindrical direction that was in the neighborhood of two hundred millionths of an inch. And how to manufacture something with a two hundred millionths of an inch bump in it looked pretty tricky to us. It turned out that there was a simple solution. That was really when we went to see the optician, and it was when we came back and figured out our own solution. The simple solution was to bend the device, the head, and then lap it flat and then let go. And that gave us a nice curvature on the bottom. And we were able then to establish parameters of load carrying, different amounts of curvature, what the angle was that it would face into the wind of the disk going by, and get it set up very nicely so that those --.

Now at the same time while the fellow I mentioned that I went to college with -- his name was Russ Bruner -- same time Russ and I were working on this we had -- a fellow named Bill Gross was doing a theoretical analysis of these bearings, and a fellow named Bill Michael was programming the theoretical work that Bill Gross was doing on a 650, an IBM 650 computer. That was about as good as we had in those days. It turned out that our experimental work led that theoretical work a little bit. Not much. And in -- I believe it was 1959 -- we published three papers in the IBM Journal of Research and Development, which are in some archive somewhere at the present time. We returned from -- after getting to that point I returned to Research, and then the Research Division -- a lot of these names come up all the time -- I worked for Al Hoagland, and we set out under Al's direction -- a person named Glenn Bacon, a person named Ray Kirby and a Jack Hildebrand -- four of us set out to build a single disk file with removable media.

Now this file was kind of unique in the sense that it had a servo system. The servo system error pattern was etched into the disk and then filled with the oxide coating in the holes of the pattern. It had vertical recording, it was a steel disk. I don't know why you'd want to have it removable because you could barely carry it, but that's what we had. And we had an interesting servo system. The servo system was a hydraulic system that was always pushing the heads either right or left. It never stopped. And you kept it on track by modulating the time. If you pushed a little bit longer to the right than the left you'd move that way. If you pushed a little longer to -- Well, the nice thing about that thing is -- and I use the word "nice" very sarcastically -- was that it whistles like crazy. It made so much noise that you couldn't stand to be around it. However, we did get a file that ran, worked, and this again was about 1959, and so we had a vertical recording system with a servo system and removable disk. It was operational, we had it on a system so that we could actually use it in the computing system. And it was essentially set up for advanced evaluation of this kind of technology.

Porter: That's very interesting, because IBM actually was an early manufacturer of perpendicular recording, which was the vertical or perpendicular recording. So it took about 45 years after you'd done that work before it became a real product. Why not make it a real product in that era?

Haughton: Yes, that was kind of interesting because my responsibility was making the heads -- one of my responsibilities. Making heads and the system that handled it. And based on the knowledge we had at that time, I concluded that it was a lot easier to make a very small nothing than it was to make a very small something, if I can say that. Because we had a probe for recording and reading, and we thought we had to make it extremely thin, and in fact they do these days, because nowadays they have deposition, and we were making it out of a piece of metal. And trying to make something out of a piece of metal that was very, very, very, very thin was a very difficult task. That was one of the reasons we went that way.

Another reason we went that way was that -- remember I said I was in the Research Division, and the file responsibility was in the Product Development Division, and there wasn't anybody over there had any skin in it, in the Product Development Division. And you've got to have a champion to make anything happen, I think, and there just wasn't a champion there to do it. As a matter of fact, they even transferred the project when we'd completed it in Research -- transferred it to Poughkeepsie of all places, where they made the large computing systems. And I'm sure that was a move to just get it out of their hair, but that's what happened at that point in time. OK, after the completion of that, the IBM 1301 was in development in the Product Laboratory, and Al Shugart asked me if I wouldn't come over and join that project in the Product Division, instead of in the Research Division. And I went over there to manage the head development and access mechanism development -- not the actuator, but all the hardware that loaded and unloaded the heads and things like that. And that was a very interesting program, the 1301. I was on that program for about two years. We did develop the heads, we did develop the access mechanism. Everything worked very well.

By the way, we did the exact opposite in getting the curves on the heads that I mentioned earlier. I mentioned that we got the curved surface by bending the heads and lapping them flat. Well, in that case we designed the lapping plates and bent them to get the curves on the heads, when we went into Manufacturing to do that. It's also interesting to say that if you asked somebody in Manufacturing to manufacture something so that it's within twenty millionths of an inch, they sort of get a little schizzy about that. And so I used to joke. I'd say "Well, we'll solve that by lying." What we did is tell them how many fringe lines on an optical block it was, and that was, like, forty, and so that was no problem. They could handle forty, plus or minus two, without too much trouble. So they manufactured them very, very well. One little thing that was of -- one problem that occurred with the 1301 was that it had a fixed track length, and you'd put so many bytes in that track, so many bits of information, and sometimes they'd come up with one more byte, and sometimes they'd come up with one less byte. Just a little bit different, one from the other. And we were having a lot of finger-pointing trying to figure out why that happened. Everybody said "It must be your circuits, it must be your this, or your that." Now the actuator on the 1301 was hydraulic. It had two stages. It had what we called a Glob Adder. You put a fixed data into a -- a fixed amount of fluid into a piston. And if you knew the size of that, why you'd move so many tracks. And it also had, in addition to that, a Piston Adder, which had finite increments within that Adder, and it wasn't dependent upon the fluid quantity, but rather dependent upon the physical dimensions of the pistons.

In the process of trying to figure out what was going on, why we had that -- we found a wonderful piece of equipment we got from Ampex. And it was essentially a frequency analyzer, so we could watch the output of these heads, and access from one point to another. And we discovered, lo and behold, that whenever you accessed, when you stopped, it would cause a vibration along the track, because they weren't perfectly centered. So the head could go upstream or downstream, and you could actually see that Doppler Effect. You could measure the frequency difference between -- when you read it back, between what it ought to be and what you were getting at that place. So we actually found that vibration in that manner, and were able -- because of the fixed glob claw -- it was cut in off -- by it or not, and we were able to solve it from that point of view. The 1301 was really the successor to the 305 RAMAC, although there were a few versions of the 305, and the 1301 was the first one to have the hydrodynamic air bearings. Even during that development, people weren't too confident. They had a little program breathing down my neck with hydrostatic air-bearing heads, just in case I didn't make it. That was an interesting competition, as it turns out.

Porter: Well, the 1301, of course, was the first production drive which used flying heads, a practice which was followed by every disk drive thereafter, from then until now, and certainly one of the most significant things that occurred in the development of the industry over those years. And it was still on those old big 24-inch disks, wasn't it?

Haughton: The 1301 had 24-inch disks. That's right. It had fifty of them. Actually, it was in two modules, at 25 apiece, one over the other, and there was an actuator on each 25 disk, but…

Porter: So there was a head flying on every disk.

Haughton: There was a head on every side of every disk. So there was a big, like a comb of heads stuck in between the different disks. And, of course, that was the part of what we were designing in my group. Of course, we had to have a loading mechanism, and we loaded and unloaded the heads, because with those heads the load on the head to hold it next to the disk was 350 grams. That's non trivial load. It's the better part of a pound. I think a pound is 454 grams. So they were carrying a pretty big load without hitting those disks. And we had to apply that load, and we applied it by an arm that came in and had two fingers sticking out, and you tipped that, and it pushed one down on one disk and pushed one up on the other disk, the way it worked. The 1301 -- you're right, it did start -- it started a trend down that direction, and all the devices after that have been with the hydrodynamic heads.

The next thing -- after the 1301 program, I was granted a scholarship to go to the University of California to work on my PhD. And they tolerated me, and I got through and got the PhD, and when I came back I worked on the Photo-Digital Store. The Photo-Digital was a photo storage of digital data, and it was a trillion-bit file. Now nowadays that doesn't sound like much, but in those days, that was big stuff. That was very, very large storage. It was done under contract to the Atomic Energy Commission and the University of California for the Lawrence laboratory in Livermore and for the Lawrence Laboratory in Berkeley. We delivered both of those the Photo-Digital. The Photo-Digital had -- the recording was done with an electron beam recorder. That was -- my responsibility was the electron beam recorder and the film processor to develop the film after it had been exposed by the electron beam. That was a mechanical engineer's Mecca -- I think it's the best word. There was more mechanical stuff in that machine -- in those machines -- I should've said it that way because the file -- the storage of the chips of film were about -- they were in little boxes about three and half inches long and an inch and a quarter square, and I think it was twenty pieces of film in that box. I don't remember for sure, but it was some number near that. And they were moved around by a pneumatic tube like they used to have for cashiers in a department store.

If you remember, many years ago you could go into Penney's store, J.C. Penney's store, and you had a little pneumatic tube, and they'd send the sale up to the accountants up in the ceiling somewhere, and then they'd send the receipt back. And that's the way these were transported, by air in these tubes. The file itself -- the storage of the little boxes I mentioned, was a set of drawers, and you could pull the drawers back and forth to select what role you wanted, and select by other means in the other direction. So it was an X-Y selection. They were transported by this pneumatic tube to a reader, and the reader was a flying spot scanner. It was a CRTA with a spot on it, and the data was coded so that you could tell when you got off track, and servo on that track quite nicely. I learned a fancy word on that program: boustrophedonic . Boustrophedonic means you go down a hill like a skier, back and forth, and that's the way our data was recorded. We recorded across the film and back, and across and back, and they were in pair -- there were black spaces and white spaces between each one for this servo system I mentioned.

Those systems were delivered to -- the first one, as I recall, went to Berkeley -- there were more -- a large lab in Berkeley. And the second one went to the one in Livermore. The one in Berkeley did not have a trillion bits. It had less, and it was used for data from their bubble chambers. And the one out in Livermore, they were full size, and used for all kinds of -- I'm not even sure exactly what all the exper… They probably didn't even tell me what all the experiments were at that time. I don't remember specifically. It was a fun project. Some people sarcastically say we proved with that project you can sell seven of anything.

Porter: Was that the total sales?

Haughton: That was the total sales, seven. I've forgotten how many million dollars apiece they were, but they were big, complex systems. After the completion of that program I -- by the way, that program was at 99 Notre Dame in San Jose, so I was off-site again during that time.

Porter: That's where the original disk drive was developed.

Haughton: Yes, that's right. That's right. That's where we did the gliding head, hydrodynamic head work I mentioned earlier -- we did that at 99 Notre Dame, and at 317 Julian Street. 317 Julian Street was a -- I'm sort of regressing here, but I will for a second -- was sort of an old warehouse. It's the only place I had to put a tarp over my desk at night because of the pigeons in the building, but it was a fun place, anyway. After the Photo-Digital work, I worked on a -- had a technology group in the San Jose lab developing advanced servo systems, advanced data recovery systems, things like that. And shortly, there became a bit of pressure from within the company that all our disk files were getting bigger and bigger and bigger, and we needed something for the low end. And…

Porter: Hence the name of the "low cost file".

Haughton: That's right. That's right. Low cost, we didn't know what low cost is, but that's what we thought it was. So there was a press -- excuse me, for a low cost file, and there were numerous people around the San Jose lab that were configuring possible solutions to that. And I remember one day that I ran into one of those people, and I said "What are you doing?" and he said "Well, I'm working on making a half-size 3330, and we'll just do that by taking out some of the disks and heads." I said "My gosh, that's a crazy thing to do." And he said "If you're so damn smart, why don't you figure out what we ought to be doing?" And that happened to be, I believe, 1969, maybe '68, I'm not sure. The day it happened was -- the day we got together on it was the day that -- President Eisenhower's funeral.

Porter: The day was March 31st, 1964. The day that IBM was closed for Dwight Eisenhower's funeral.

Haughton: That's right. That's right. And -- not 1964 was it? That doesn't sound right, because we didn't -- we shipped the Photo-Digital in 1968. Well, anyway, it happens the meeting occurred on the day that Dwight Eisenhower's funeral was, and they'd shut the plant down for that day, and about four of us came in and did a little brainstorming, you know, on concepts. And that's when we decided that we ought to move on to much smaller heads, much lower mass, much lower loads, so that we wouldn't have a risk with the disks as well as- with that 350 gram load I mentioned a while ago, although that had gone down in the 2314 and 3330, along in there.

Porter: So the four of you did agree on some kind of a concept then?

Haughton: We agreed to start out with that we would first of all go to small heads with a small load, and we'd try to land and take off on the disk, instead of having a loading mechanism. We, in fact, concluded that if, that were the case that it might be that the best thing to do was to have the heads and the disks matched up because -- in a set -- because that would eliminate a lot of manufacturing tolerance problems. Wherever the head is, that's what it records, but if you have to read it with another one, then you've got another problem -- manufacturing tolerance problem, putting it together. So we agreed that we should set out to try and put the heads right with the disk. As a result of that, not too long later as a matter of fact, I was asked by the Product Manager to come away from my Advanced Development Group and lead the development of the Winchester Program.

Porter: And obviously, we need to ask you how you came up with the name Winchester, since all IBM projects had to have a code name, as I recall.

Haughton: Yes, all IBM projects had a code name, and some were kid of humorous, if you knew the history behind them. For example, the 3330's code name was Merlin, and they called it Merlin because you had to be a magician to make it work. <laughter> And in the case of the Winchester Program, the title... That came up because our initial configuration was a box with two spindles in it, and the two spindles each had thirty megabytes of storage. And we found people calling it a 30-30, and I said "Well, if it's a 30-30, it must be a Winchester," and that's how the name got stuck on it. And by the way, it stuck with that technology for many years. It was a long time later before it finally began to just be "hard disk".

Porter: Oh, many years it was normal for people to talk about the disk drives being developed in the seventies and eighties as Winchester Technology.

Haughton: Yes, yes. Whoever was doing it, it turns out. As a matter of fact, it was the basis for a big expansion in the industry at that time. But anyway, we put this program together, and I transferred to the Device Group, if I can call it that, within the Product Development Group, and led that development. It took us about four years. I pride myself in one thing, and that is that I made a schedule for a ship in the fall of 1973, and that's when we shipped. Shipped in November of 1973. Maybe that's late fall, but we came pretty darn close to the schedule that we set out in the first place. The data module, as I mentioned, was kind of interesting. We had a -- did develop a scheme to include the heads and the disks in one package, separate from the drive, and those -- that data module, of course, had a lot more technology in it than the disk packs of the previous devices, because of the heads and what have you.

We went -- we had a case, a dispute going on in the company as to who would sell disk packs during the development of the Winchester Program, and, based on the 2311 and 2314 disk packs and so on, the corporate decision was to have the Supplies Division sell the disk packs. And then they realized, here was the Winchester file- project coming along with what we called a data module, which had not only disks in it, but heads and a carriage, and the whole business, so they decided "Well, we'll just let the Supplies Division sell that too." That was what I call a million-dollar -- a million-mile memo -- when I got that memo saying it was going to be sold that way, because I could visualize all kinds of problems in the field, if somebody selling supplies is trying to sell a technical piece of equipment. So I spent about a million miles on airplanes going around arguing “What can we do to keep this from bombing out for the wrong reasons?” And it turned out we ended up with a compromise. We ended up incentiving the sales force by giving both the data processing salesmen and the supply salesmen a commission for selling those data modules. However, it wasn’t very well communicated to the field, and as I went out to visit the field, the sales officers and customers -- I’d find them selling what I thought was the wrong device for the application. And I said “Well why don’t you use the Winchester one instead of the 3330,” if that were the case. And the guy would say “Well, because I don’t get the commission.” I said “Well, wait a minute, now, I spent a million miles on airplanes making sure you would get the commission.” So eventually we got the word around, and they did sell those, but it was an interesting communication issue that -- a tension that a person doesn’t necessarily think about when they’re working on a product development line.

Porter: Let’s hold it here, and let them change the film.

Haughton: Oh, OK, we’re at that part.

[RECORDING TAPE CHANGED]

Porter: Well, Ken, the Winchester program, in addition to being hard to sell to some of the IBM sales force, was, I'm sure, one of the harder product development programs that ever got put together. It's widely known that when it's finished, this is the first program that really combined what was known as a low mass head, a lubricated disk, a sealed enclosure. And, of course, those three principal characteristics then ended up being used on all disk drives since the Winchester. But what were the more important and difficult challenges to getting all that technology put together?

Haughton: Well, you know, it was not easy. We had close to 100 people in the machine group developing the machine itself, and that involved the data module, how it all went together. It involved the mainframe -- if I can call it that -- the box where you mounted the data module, the power supplies, and all the standard equipment. We had close to 100 people working on that. In addition to that, we had another large group -- I won't say it was 100, but it was a large group -- developing, 'cause we foresaw the need for real mass manufacturing of those heads. So we had a large group working on the head, both design, manufacturing... Working together with manufacturing they had the head group going. And then we had a group of a dozen or so working on read/write and servo and that kind of thing. So all of those things were consuming, as far as effort was concerned. And as far as the major, we didn't have the problem like we did back in 1301, where, "Can we keep these things going or not?" we didn't have that problem. The heads were fine. We still had to work on them and we had little nits, but there wasn't a counter feeling going on that you might not be able to do it with those heads. Everyone was on the believer wagon by that time. But there was still a lot of work to be done, and I can't think of anything specific, other than cost, that really laid in the way. We were trying to get a low end file -- remember, I said when we started this -- and trying to get the cost down to a reasonable figure.

Porter: Well, Ken, those heads were smaller than any heads previously used, so low mass. There was some derivation from some of the activity IBM had been involved with, with a license they purchased from Data Disk on their small heads, and there was some kind of a melding of the work that had been done on the IBM 2305 drive in all of this to pull all this together, wasn't there?

Haughton: Oh yeah. Well, first of all, let me address the Data Disk heads. IBM did license the Data Disk head, and Joe Ma built a video file in Advanced System Development division using that head and 24-inch disks. Did this under contract with the Rand Corporation, and it used those heads. That's where we started with the Winchester as well. These heads, it turns out that mechanically, they were not as stable as we had liked and a bit difficult to manufacture, because there are three separate little pads, air bearing pads, on any one head. And any variation in any one of those could affect the way it flew. The fellow that had charge of the development of the heads, a fellow named Eric Solyst, had developed the heads for the IBM 2305. And the IBM 2305 had a head that was what we called a tapered land, as opposed to the curves of the cylindrical ones that I mentioned earlier. And he was like anybody else, wanted to go the way he knew the most. So he drove this toward making a tapered land. Now, it turns out the tapered land has some real advantages from a manufacturing point of view, because everything is straight cuts. So you can grind, you can lap without any tricks like bending the lapping pad that I mentioned previously, you can do it. So that was a big argument for doing it that way. And it turned out that that's the way we went, was with the tapered land. I had some, I guess, pride of ownership or something in the cylindrical ones because I was interested in going that way. Even this day, I have some reluctance about the way we did it, because this introduced the stiction problem in the business, having a lubricant on the disk surface. And the disk surface and the head surface, when they're stopped being absolutely parallel, could stick together like nobody would believe. And sometimes we'd had situations bad enough that we'd rip the heads right off when we start up. So it created some new problems. But in the long run, the manufacturing thing was good, and who knows whether we would have solved doing the curved one well or not.

Porter: The technology you used on those heads and the disk surface, when the 3340 was finally introduced using it, the head-to-disk flying height was 18 microinches. Which was, by far, the narrowest distance between the head and the disk surface that had come before. Before that, the other drive introduced in the seventies, a 3330, flew at 50 microinches. So this was a very significant advance in getting that head closer so you could record at greater densities and so forth, wasn't it?

Haughton: That's right. It was a big chance -- and if you look at what's happened now, it's equally, and in fact, perhaps even greater. They're getting down in the one and two microinch range, or somewhere around that nowadays. But that was a big change. If we go back to the 1301 when we started this business, that spacing on the 1301 -- I just scratched some cobwebs out of my head here --

Porter: Eight hundred microinches.

Haughton: No, no, that's the RAMAC.

Porter: Yeah. Excuse me. The 1301 was 250.

Haughton: Yeah. The RAMAC was 800, and so we came from 800 to 250 to, I don't remember what the 2314 was.

Porter: 2314 was 85 microinches.

Haughton: Okay, so you've got the data there. So, anyway, spacing is the big key on recording density. It really makes a big difference. It, first of all, allows you to get higher bit density, and second of all, allows you get higher track density. So both of those things contribute to the storage capability. Nowadays, I don't even know what the track density is, I haven't looked lately. But it's in the thousands somewhere. Thousands of thousands I think, almost. It's really gone down. Or up, whichever way you want to call that. Much more dense. Let's see, where were we?

Porter: Well, the heads are developed and the disk -- you apparently tried various methods of lubricating that disk, didn't you?

Haughton: Yeah. I mentioned a minute ago that Joe Ma had done this video disk with the data disk heads. And he introduced the concept of lubricating the disks to improve the start/stop characteristics and any damage when a head accidentally hit the disk or something like that. And so we did that on the 3340, the Winchester, and in fact, it became standard on all the disk drives through that time frame. And people still do it. The lubricant has changed, but they're still doing it. It's a different lubricant.

Porter: But you had to try many different choices on that lubricant, I guess, didn't you?

Haughton: Well, yeah. During that time frame, there was a big hullabaloo about what's patentable and what's secret, and what's the difference. And the lubricant got classed as a trade secret. And so there were very few people that knew what lubricants you're using. It was sort of a cloak and dagger type thing, to keep that so confidential, because it was pretty key on whether the devices worked or not. And no one could see any way to patent it, and so we chose to use the trade secret route with the lubricant. So there were different ones. If you remind me, if we get to a little bit of the 3380 story, I'll come back to that because it was a key item in the 3380. But the Winchester drive, as you say, set a trend, and we have -- you know better than I how many companies made disk drives in that time frame, too many to count. Now it's getting back down to a rational number again, but the whole business really blossomed as a result of the Winchester.

Porter: There was another interesting touch where you used two heads for a disk?

Haughton: Two heads on a side. I hadn't even thought about that for quite a long time. We did have two heads on a side. If I go back to the single disk file that I mentioned that we did in research, we had four heads on a side on that one. And this tradeoff on how far you move the actuator, how much area is covered by those heads, and how fast you want to go. And we elected in this particular case to move the actuator half as far as all the previous disk drives and to put two heads on each side. I'm not aware of any since then that did that. There may be some, but I don't recall any.

Porter: Well, the heads were small, low mass. You could get to the data faster that way.

Haughton: And we were getting the cost of the heads down. I mean, the cost of the heads on a 1301 compared to the cost of the heads on the Winchester were enormous. So we saw that as a way to improve performance without pushing the actuators and what have you.

Porter: The bottom line was that you got a much higher tracks per inch number, the 5,636 tracks per inch.

Haughton: Oh, that's the --

Porter: The TPI, excuse me. And then on tracks per inch, you got up to 300 tracks per inch, which was way up from the 3330's 192.

Haughton: You just reminded me of another story. I had a weekly staff meeting of all my managers and all the managers from these other groups that I mentioned that were doing the read/write stuff, and the disk stuff, and the head stuff, and so on. And we started the Winchester program when it was 30 megabytes per spindle and 167 tracks per inch. But then as I traveled around with the DP marketing folks and what have you, I began to realize we weren't very competitive. The device was not competitive and we had to do something. So at one of these staff meetings, I said, "I want each one of you guys to tell me what you can do to increase that density." Tell me, the read/write guys, "I'd like to go up in bid density and up in track density." The head guys, "I need a head that'll do that." The disk guys, "I need a head..." And so then, after whatever investigation they did, why we got together and at the staff meeting, I said, "Well, let's poll everybody now," and I polled each person. And I said, "What can you do about the read/write circuitry?" And the guy says, "Well, I can go up quite a lot." I said, "Well, what about 200, 300?" The disk guy said, "Well, I can do a lot better than I'm doing now, except I don't think they can make a head." And then I turned to the head guy and say, "How about you? Can you make a head that'll do this?" He says, "Oh yeah, I think we can do that, we can get there. But I'm worried about the disk." And so on, all the way around the room, everyone worried about somebody else's territory. So I arbitrarily picked 300. I've kicked myself ever since, I didn't pick 500. I think they would have done it, as history now tells us that it was there. Whether we were ready for it or not is the question.

Porter: So that's amazing. The 300 tracks per inch was a target you picked in advance, and they actually got there.

Haughton: Yeah. And I just did it in a meeting like I just discussed. It was, "That's what it looks like we got to have in order to have a competitive device," and it was the same kind of discussion about the bit density too. I notice that you have on your notes, when I was looking at them last night, that it goes some number like 2965 RPM, 64, whatever it is. And that's true, that's what it was. But that's because it's an induction motor. We set out for 3,000 and the motor just slips a little, it goes a little bit less than that.

Porter: Pulling all this together was quite a management challenge, and of course, you had managed several projects before. But when you took all of this over, and you had the task of making all of these various diverse areas of the technology work together and somehow hit those targets, how did you feel about the management challenges you were facing?

Haughton: Well, that's kind of interesting, Jim, because at first it was chaos. Getting it communicated so that everybody bought in was not well done. And finally, I decided to take these same key people that I just mentioned a minute ago -- primarily, that was about 20 or 25 people -- went down to the La Playa Hotel in Carmel, and rented their big room they have up there. And just had a -- we used to call them "love ins". But what we had was a discussion of what we're trying to do and why. And it really showed how poor the communication had been. It really showed what needed to be done from that point of view, and I spent probably four or five hours, in a sense just giving a history lecture on how we got to where we were. And that pulled the team together. Everybody understood the history and where we'd been, and where we needed to go, and they started pulling together very, very well. If we hadn't done that, we probably wouldn't have made it. I think it was that significant of a session.

Porter: Was it a very big challenge after that agreement in principle on those things?

Haughton: Oh, yeah.

Porter: You had reached agreement on all of those things. How much of a challenge was it to really get the final finished product done?

Haughton: I think as far as getting a running machine, that was fairly smooth sailing. Probably an overstatement, but we got along pretty well. The problem we had was it cost too damn much. And we actually brought a guy named Joe Sheredy over from Manufacturing, and we just pounded that thing and pounded it and pounded it to get the cost out of it. And the cost was a very serious issue, probably the largest jeopardy, because, again, we started out for low-end systems. And they were the ones that were going to use it, and they couldn't afford to have too much invested in data storage. So the cost was the biggest point from that point on, I think, although there was a lot of work to do. IBM used to have an organization -- I don't know whether they still do or not -- where you'd turn the machine over to a product test and they'd test it and kick the covers and everything you can think of. And they'd do a good job of demoralizing you, because the first problem list has got like 300 things on it. And you'd say, "How can there possibly be 300 problems here?" But there is, because they cover everything, right down to "the hinges don't work on the covers," or something like that. So there's a lot of problems solved.

Porter: Now, the original product was going to be 30 megabytes fixed, 30 removable, on different spindles. Of course, the final product was data modules of 35 megabytes and then 70. And then, there was a third version, as I recall, of 70 megabytes with some fixed head options on it. And then, of course, you had a fixed on the 3342, didn't you?

Haughton: No, we didn't have a fixed.

Porter: You didn't have fixed. It was all removable, the final one.

Haughton: Yeah. They were all removable, and they were right from the start. I will say this, that part of that motivation behind the data module concept was the vision that we would return to fixed files. And the reason we expected to return to fixed files was that by having the heads to the disks as paired couples, you can drive the manufacturing costs to where you can afford to have fixed files as opposed to removable. Because that was the problem they were solving when they invented the disk packs in the first place, was to get the system to have more data available without spending a fortune for a whole 'nother machine. And so we had that in the back of our minds -- really, not in the back of our minds, it was in presentations and everything else right along -- is, "We're going to head that way in the long run." We did some surveys on how much they changed packs on 2314, and it was not much, on the average. You cited to me earlier today some cases where they did it a lot, but on the average, the number was not very much. And in fact, when we started pushing the fixed files, the marketing people would say, "Well, how come we've got to have a data pack, a disk pack?" And I said, "Well, what do you want those for?" They said, "So we can take them off." And it turned out that the concept had become so embedded in people's minds that they thought they had to take them off. But it turns out, as you know now from history, that in fact, the prices have gotten such that taking them off is not worth it. So we did go that route.

Porter: As you finished with the Winchester project and it was introduced as the 3340, you had managed through some difficult times and with a lot of interesting management techniques to pull it all together. You would manage successfully the project which produced probably the most widely known term and name in the storage industry, Winchester. And you had managed the project which had, perhaps, the greatest influence in the future on products that had come at any time. How did you feel about having been involved with all of that, in that role?

Haughton: Well, you know... Very good. As a matter of fact, I had an interesting thing happen at the end of that. IBM had, at that time -- I can't attest for them now, since I've been retired from there for 25 years, it's kind of hard for me to know what they have now -- but they had numerous awards. And those awards would vary all the way from dinner for two to a large cash grant. In 1964 I bought myself a four door Dodge pickup and took my family to Alaska and did things like that. I got a camper to put in the back of that pickup so we could all ride in the cab and so on, and I nearly wore it out. And as the Winchester program came to an end, one day I was talking to a staff guy and he said to me -- it was an appropriate buildup, so I didn't suspect a thing -- he said, "How do you like your pickup?" And I said, "I love it. It's really great. I've been able to do this, and this, and this and so on, I love it." He said, "Well, what would you do different if you were to buy another one now?" And I said, "Oh, I don't know. I have pretty much what I like." And he said, "Well, what about four wheel drive?" I said, "No, I don't think I'd want four wheel drive. I might get a winch on the front, but I can't think of much improvement." And he said, "Oh, okay." And then we went on with whatever we were doing, and that was it. And a few weeks later, I saw him. He said, "Come here." And we went outside and there's a brand new pickup. And that was an interesting mistake they made during that. In order to buy that pickup, they had to treat it like a cash award. And so they found out how much it cost and got that much cash. Well, it turns out that the payroll department took tax out. So they didn't have enough money to buy the pickup. So they went to my wife, and she dug up the difference. So anyway, I used to have a -- I think I still have it somewhere -- a little brass plate that I had on the dashboard saying, "Recognition of doing these things on the Winchester program," signed by Art Anderson, who at that time... I can't remember. I think he had already gone from division president to group executive. But anyway, he was somewhere in there.

Porter: That's a unique recognition of achievement, to give the manager a pickup truck.

Haughton: As a matter of fact, our personnel people just ate it up. They'd use it when they went to interview company personnel things and say, "Well, we gave this guy a pickup," and that blew everybody's mind. It was a very unique and interesting thing.

Porter: So today, all these years later, when you look back on it: Still satisfied with having done all of that, I would imagine?

Haughton: Oh, it was a fun thing, yeah. It was a very fun thing, and a real achievement. I think a major achievement. I feel very good about it.

Porter: Excellent. Then completion of the program. You moved on still with IBM?

Haughton: Yeah. The completion of the program, I went back into the laboratory -- I wasn't out of the laboratory in the program -- but I took a job managing the -- I'm trying to remember the sequence right here.

Porter: Well, wasn't there this program called Future Systems?

Haughton: Yeah, that's where I was headed. At that time, IBM was developing what they called Future Systems. And System 370 was sort of a "You bet your company," and this was sort of a "You bet your company." Very major project, very big all over. And I was asked to manage the storage program for that. And I had the storage system, which had people both in San Jose and Boulder. 'Cause Boulder had tape devices, and we were working on staging data from one kind of memory to a faster, to a faster, and to try and predict what data you'd need and that. So we were working on putting that together for that Future Systems. That program within IBM went down in flames, if I can say that. They never did finish it, as they did the System 360 and 370. So I was on that for a couple years, and then I was asked to come to manage the ASDD laboratory in Los Gatos, Advanced Systems Development Division. I went to that job in 1974, and that location didn't do very much, in the way of disk files.

I mentioned the video system that Joe Ma did earlier. That was in that laboratory. There were a few other magnetic recording things going on that he and Dr. Emil Hopner... But mostly they were working on displays. The IBM ATM, automatic teller machine came out of that laboratory. They did the BART system, the BART ticketing system. So I went over there to manage that lab. Its five miles down the road or something like that. And after about a year, we concluded that -- the size of the lab was fairly small, about 400 people -- and we decided they would be more effective if it was merged with the San Jose lab. So we did that, and I went back to manage the technology in the San Jose lab. Considerably bigger than it was when I did that before, but same concept. I was there for -- I don't remember how long it was. But then, in 1977, I was asked to go to Lexington, Kentucky and manage the typewriter printer facilities in Lexington. By the way, in 1977, they already had out in the world an inkjet printer, which people don't really realize very much, I think.

And I have a little side story about an inkjet printer that I'll give you at this point. That particular printer, the ink was flowing all the time. Drop after drop after drop. And it had an electrostatic deflection system to deflect those drops you want so that they would print on the paper, and the rest of them went into a gutter and just recirculated. There was a problem in inkjet printing, at least with a liquid ink like that, that there's a trade off between how well it adheres to the paper and how sharp the image is. Well, one morning, we were having trouble with the ink in the field, and so we were trying to develop a new ink at the time. And one morning I got up, and before I went to work, I read the Wall Street Journal. And I read an article that said some folks had done some tests with the dye DB-19, and found that if you feed big doses of it to rats, they get cancer and die. And I thought, "God, that rings a bell." So I went in, and sure enough, DB-19 was the dye we'd used for the ink that was out in the field. And, of course, in that time frame there was a big panic about any carcinogens, there was news all the time. So I said, "Well, what are we going to do about that? We've got this other ink that is not quite ready, doesn't use DB-19, but we've got a bunch of it out there. So we decided to put a label on the ink. And the label said, "Do not drink this ink." And we got it translated into eight languages, and labeled all the ink and sent it out. And one day a couple months later I got a phone call, and the guy called me from Seattle. And he said, "You're the outfit that sends this ink out, aren't you?" I said, "Yeah." He says, "It's got a label on it, it says, 'Don't drink this ink.'" I said, "Yeah, you got the right one." He said, "Well, in Spanish, it says, 'Drink this ink.'" Fortunately, we were ready to ship the new ink by then. So you get little stories that happen in all kinds of places.

After leaving Lexington, I came back to San Jose as the manager of the San Jose laboratory, director of the San Jose laboratory. I got here in September of 1980. And the 3380 was in CTEST. CTEST, in IBM parlance, is the test of the manufacturing ability of all the documentation and stuff you have released. So it was ready, it was time to ship, just about. And they started breaking, one right after another. Just bang, they'd crash. Bang. One after another, they crashed. So I had a crisis on my hands as quick as I showed up, and that's what I said I'd come back to. It turned out the problem was the lubricant on the disk. They had found a lubricant that was much better, from a point of view of lubricating the thing, from the point of view of stiction, a whole bunch of things. It was much better than the one they'd been using, so they switched. And it took quite a bit of work to figure out that it was the lubricant, but in fact, we did. And the following would be my...

[RECORDING TAPE CHANGED]

Porter: You can just pick up from the beginning of that last part.

Haughton: Okay, well, I'll pick up that I came back to San Jose as a Laboratory Manager in September of 1980, and the 3380 was in C-Test, which is the manufacturing ability test, and they had been running for about, oh, for several months, and suddenly they started breaking one after another. By that, I mean head crashes, and that turned out, after a lot of busyness and a lot of work and so on, to be lubricant. They'd found a lubricant to put on the disk that was all apparent characteristics, it was better for stiction, better for start-stop consequently, and almost anything else. It just was a better lubricant than the one we'd had before, so they made this decision to change sometime in the past. Well, the rest -- so we changed back to the other lubricant, and got back on track, and, of course, during that process, and going back to IBM and telling the headquarters we weren't going to ship that baby the way it was, which is always fun to do that.

But anyway, I have a personal view of that, and this is personal. It turns out that the, I'll call it the disintegration temperature of the new lubricant was a little bit lower of the old lubricant, and I did some calculations over the years that followed by just calculating the energy of sheering off any peaks that stick up, and lo and behold, I could predict that temperature that we would reach in the first one, and not the second one. So, in my view, what was really happening was we had a very good lubricant, but it was too easy to heat it all up, and so it was gone after some period of time running, because you do have some intermittent contact on little peaks on the disk that stick up. After we did get it back together, did get it running, and as you know, it was one of the most successful programs that IBM ever had...

Porter: After an embarrassing several months of delay, though, right?

Haughton: Yeah, embarrassing. I was the embarrassed guy there.

Porter: I think the folks that really enjoyed the change were the folks in Colorado at Storage Technology, who shipped a whole lot more 3350's than they knew what to do with.

Haughton: Yeah, that's right. They got the 3350's out like mad.

Porter: You sell a capacity version, and they shipped them like crazy, yeah.

Haughton: Yeah. Yeah. So about that time, I got a call from Jack Kuehler. Jack Kuehler said to me, "Did you mean that?" and I said, "What are you talking about?" He said, "Well, when you went back to the university to get your PhD, I asked you one day if you thought you might like to go back to a university some time," and I said, "Well, maybe. I don't know. I hadn't thought about that. But I'd consider it." And that had been 20 years before. And he said, "Well, the reason I'm calling you is we're looking for somebody for Dean at Santa Clara University." And the upshot of it is I ended up taking that, resigning from IBM and taking that job at Santa Clara University.

Porter: And you were there for quite a few years.

Haughton: I was there from February of 1982 to the middle of 1989. I tell people I retired three times. I retired from IBM in 1982. I retired from Santa Clara in 1989, and while I was there I was on the Board of DaVinci Graphics, and they were having some troubles and, so after I retired at Santa Clara, a year or so after, I went in and to run their engineering effort for a while, and then retired from that. It finally took.

Porter: Well, you've may be retired, but still pretty busy. I guess I know you were on a couple of other boards, such as Solectron, and Seagate, of all things.

Haughton: Yeah, I was on Solectron and Seagate Boards, each of them for over ten years, and about four other companies. DaVinci, I just mentioned, I was on their Board. I was on the Board of Viacom Systems, which was an outfit that made medical imaging equipment. I was on the Board of a little outfit that made medical chemical, or chemical instruments, I should say. I can leave the medical out of that, although they were for medicine, and then I was on the Board of, I can't think of the name of that company or the other one. The other one still exists. It's down in Tucson, and makes optical equipment, fancy optical equipment. So I had about four other companies, besides the two you mentioned. They were all private, and in fact Solectron was private when I went on it, but it went public during the time that I was on the board.

Porter: In the case of Seagate, there were some people there that you'd worked with at IBM in the old days.

Haughton: That's right.

Porter: They're probably the ones that recruited you, no?

Haughton: Yeah, it is was Al Shugart that called me, and asked me if I'd be interested in it, and there were several IBM folks around the place there. He always attracted this crowd that moved with him wherever he was.

Porter: Al?

Haughton: Al did, yeah. A fantastic guy, and so he attracted people who liked to work with him.

Porter: And you went through some very interesting and exciting times with Seagate?

Haughton: Yeah. Yeah, they absorbed a couple of companies around, and eventually parted company with Al, which was no fun. But it was, still is its leader in that field nowadays.

Porter: Well, they're the leader in quantities and revenues, yes.

Haughton: Yeah. Yeah. They are a quite successful company. I can't say that I had too much to do with that.

Porter: Well, you had a chance on the Board to observe and certainly call shots if they're not doing things properly.

Haughton: Oh, that's right, and we had some times when we really got pretty close. The Board got pretty darn close two or three times where some of us ended up essentially down there a couple days a week.

Porter: You were at Seagate when they forced out Al Shugart, weren't you?

Haughton: Yeah, yeah.

Porter: Interesting times.

Haughton: Yeah. It was a very interesting time, and as I say, that was not too much fun. Although I have to say this, you know, I think it ended up exactly the way he wanted it. I think that he was tired of that job, and wanted to move on, and so I could be wrong, but that's my feeling. I feel like that's the way it was.

Porter: Well, prior to that I had recruited Al to be the keynote speaker at the Diskcon tradeshow, sponsored by IDEMA that year, and I happened to go into the office early the day that they put out the press release, and I noticed at 7 a.m. on my fax this thing on Al leaving, and so forth, at the request of the board. I went in my office, I picked up my phone and called Al on his private line. He answered, and I said, "Al, I assume you're kind of busy this morning, but I just wanted you to know that when I asked you to be the keynote speaker at Diskcon, I invited you, not the company," and so Al said, "Yeah, thank you. I'll call you back in a couple of days and confirm." And he did.

Haughton: Yeah.

Porter: And he gave his speech to more than 1000 people in the ballroom of the Fairmount Hotel, and it was essentially Al's farewell speech as an officer of the company, and it was, as I remember, one of the industry's great occasions. Everyone afterwards stood up and gave him a standing ovation for as long as I can remember.

Haughton: Well, yeah, he was very good at that. Very good at, I call it impromptu speaking. I mean, he can stand up and make you live the story right with him. He was very good at it.

Porter: It was a memorable evening, and this is a, I think, a memorable occasion today. Very much appreciated, you spending the time to go over your contributions to the industry, which have been very significant to the disk drive industry.

Haughton: Well, that may be a little bit of an overstatement, but I'll accept it.

Porter: Thank you, and I think it's without doubt a truthful statement. Thank you very much.

Haughton: Thank you.

END OF INTERVIEW

